

Generating Tailored Explanations with an Independent Agent Model

*Kalina Bontcheva
University of Sheffield
Regent Court, 211 Portobello Street
Sheffield, S1 4DP*

[*K.Bontcheva@dcs.shef.ac.uk*](mailto:K.Bontcheva@dcs.shef.ac.uk)

The field of Natural Language Generation (NLG) studies the automatic production of coherent texts from some internal representation. Often these texts are made sensitive to the context and the target audience: explanations for novices contain different information from explanations for experts, for example.

A number of previous NLG systems have used different types of user information to generate user-sensitive texts. However, none of these systems have attempted to use a comprehensive user modelling framework. The most common representation is an overlay model where facts from the generator's knowledge base can be marked as known by the user. Often the representational power and contents of the user model are chosen to be sufficient for the task at hand. However, the model might not always scale up well or be portable and re-usable.

In our view, what is needed is an independent, powerful agent modelling component which can be used to support both understanding and generation. The advantage of using such a component is that potentially it could be reused across application domains and also between generators. In addition, having the agent model as an independent component could facilitate the evaluation of its influence on the various generation tasks.

In order to study these issues, we have examined an existing agent modelling framework (ViewGen) which we are now integrating into a language generation system. ViewGen is an agent modelling system which represents beliefs, intentions and goals of dialogue participants. The framework is based on nested belief environments corresponding to agents (system and user in our system). ViewGen computes these environments on demand by a process of ascription.

An important advantage of ViewGen is its account for common knowledge using default ascription and stereotypes, as this allows modelling of different classes of users and tasks and use of the most relevant one(s). Also, for efficiency reasons, the views of each agent are built on demand and relevant stereotypes are consulted with respect to a specific belief or topic only when necessary.

The current implementation runs on a knowledge base in a chemical sub-domain. A bigger experiment in a computer science sub-domain is in development. The system contains a content planner and a surface realiser; these modules are executed sequentially. To limit the task, the agent model (ViewGen) only affects the selection of relevant content and its organisation into a structured explanation. The current text planner is schema-based and we are using an existing surface realiser for English. The generated explanations are paragraph-sized and contain

hypertext links to related concepts. The initial version does not use any user model, and neither does it take into account previous utterances. The intention is to use it as a baseline system against which to compare ViewGen's influence and performance. In addition, this incremental approach prevents the generator from being dependent on the user model, i.e., it is able to produce acceptable texts even without it.

The influence of the agent model will be evaluated by comparing execution time and output quality with and without the model. A black-box evaluation measuring the user satisfaction with the explanations is also envisaged.